DOCUMENT RESUME

ED 084 809 EM 011 620

AUTHOR Hawkridge, David
TITLE Problems in Implementing Computer-Managed Learning.
INSTITUTION Open Univ., Walton, Bletchley, Bucks (England).
PUB DATE 73
NOTE 18p.

EDRS PRICE MF-$0.65 HC-$3.29
DESCRIPTORS Behavioral Objectives; *Computer Assisted Instruction; *Computers; Computer Storage Devices; Diagnostic Teaching; *Individualized Instruction; Information Storage; Man Machine Systems; Problems; State of the Art Reviews

ABSTRACT

The demand for computers to help manage learning comes mostly from those engaged in attempts to individualize education, who have discovered a need for computers to help them make decisions and organize their data. This paper considers the problems that can arise in implementing computer-managed learning by indicating what has happened in existing, operational projects. First, it describes some of the problems in diagnostic assessment of learners--the process of discovering what repertoire of behavior each learner has already acquired and where the gaps are in that repertoire. Second, the paper examines problems in the generation, analysis, and banking of test items. Next, problems in the selection of a particular sequence of objectives for an individual student by the computer are discussed, and finally, some of the problems that teachers and students experience in using the computer system are identified. (Author/SH)

ED 084209

PROBLEMS IN IMPLEMENTING COMPUTER-MANAGED LEARNING

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

Introduction

This paper is about problems in using the computer to help manage the learning process. It is not about computer-assisted instruction (as that term is generally understood), computerised class-scheduling, computerised educational accounting and financial control (Bushnell and Allen, 1967), or the linking of computers in networks (Weingarten, et al., 1973) although in all of these the computer is indeed being used to help learning.

This paper deals particularly with problems in four areas:

1) diagnostic assessment of learners
2) the generation, analysis and banking of test items
3) the sequencing of learning materials for individual students
4) the interface of students and teachers with the computer

These areas were selected at the request of Richard Hooper, Director of the National Development Programme in Computer Assisted Learning, for which the paper was prepared.

The source of demand

Where does the demand come from for computers to help manage learning? It certainly does not arise in schools or colleges in which traditional 'lock-step' methods of teaching are employed. Teachers who assume that all their students are closely similar, and that these students will proceed at the same rate by the same route through the matter to be learned, will also assume that the computer has little to offer.

The demand comes from those engaged in attempts to individualise education.

About 50 years ago, Washburne et al. (1926) and Parkhurst (19[illegible]) reported that teachers using the Winnetka Technique and the Dalton Plan, two schemes for individualised learning, found diagnostic assessment and record-keeping functions most time-consuming, both for teachers and pupils. Yet the data volume in those schemes was small by comparison

EM 011 620

FILMED FROM BEST AVAILABLE COPY

with the volume used in two large American computer-managed individualised education projects now in operation. Nor was there any attempt in those schemes to evolve sets of instructional decision-rules (to guide students) that used the data.

In the United States, new impetus was provided for the individualised education movement by Project TALENT (Flanagan, et al., 1962). Project TALENT, itself made possible by the computer, involved the collection in 1960 of nearly 2,000 items of information for each of 440,000 high school students. The data were banked in a computer, which subsequently produced a great number of analyses. The students were followed up one, five and ten years later, but it was the initial analyses that demonstrated in some detail the need for individualising education. Flanagan (1971) stated that 'between 25% and 30% of ninth grade students had already achieved as much knowledge and ability in such fields as English and Social Studies as the average twelfth grade student. This great variability in the level of achievement of the students in a particular class suggested a very real need for individualisation...'. Flanagan also found that TALENT data reflected students' feelings that their high school courses often did not meet their needs, and that there was a lack of stability and realism in the students' occupational choices at that time.

TALENT led to the foundation of Project PLAN (Program for Learning in Accordance with Needs) in 1966 by Flanagan. At about the same time Glaser and others (Cooley and Glaser, 1968; Glaser and Nitco, 1970) in Pittsburgh conceived IPI (Individually Prescribed Instruction). These two are major computer-managed individualised education projects into which many millions of dollars of development money have been poured. PLAN represents the most comprehensive use of the computer, offering complete curricula in mathematics, language, social sciences and mathematics for both primary and secondary levels. IPI is a more limited scheme, covering only mathematics, language and problem-solving techniques at the elementary level.

The Open University represents a 'utility' model in its use of the computer, and does not offer as much individualisation as the other two. From the three, a fine range of examples can be provided of the problems that arise in implementing computer-managed learning.*

Four problem areas

Computer-managed learning systems are built on a scientific, behaviourally-based model. (This could be an adverse comment!) The systems are built by people who believe it is possible and useful to analyse, categorise, measure and change the behaviour of students (see Cooley and Glaser, 1968; Anastasio and Morgan, 1972; and Glaser, 1969). Most of the problems in their systems arise in their attempts to carry out these functions. Frequently the problems are the result of sheer human diversity: there are too many variables to take into account. Even the 1,000,000,000 items in the TALENT data bank cannot provide exactly the information that is required for certain important decisions about a student's learning. Other problems arise from the need to blur discriminations, to over-generalise from specific data, and so on.

Such problems are not always solvable, but this paper at least indicates what they are like in existing, operational projects. First, it describes some of the problems in diagnostic assessment of learners, that is, the process of discovering what repertoire of behaviour each learner has already acquired and where the gaps are in that repertoire, assuming that for him certain sets of educational objectives are desirable. Second, this paper examines problems in the generation, analysis and

---

*Footnote: There are several other projects that would yield further examples of problems: e.g., the Computer Managed Learning Project at Colorado State University, the Instructional Management Systems Course at Western Washington State College (Latta and Straughan, 1972), and (possibly no longer operational) the Instructional Management System developed jointly by the Southwest Regional Laboratory for Educational Research and Development and Systems Development Corporation, for Los Angeles elementary schools (Gaddes and Kool, 1969).

banking of test items. Since test items are intended to be valid for certain educational objectives, these objectives enter the discussion again. Third, the selection of a particular sequence of objectives for an individual student by the computer contains a further set of problems. Fourth, this paper identifies some of the problems that teachers and students experience in using the computer system; these are commonly called interface problems.

Problems of diagnostic assessment

Fundamental to any computer-managed learning system that attempts to match objectives and materials to learner characteristics and needs are analyses of each of these: learner characteristics and needs, learning objectives, and learning materials. Diagnostic assessment involves identifying the learner's characteristics and needs. Its first purpose is to make possible logical selection of appropriate objectives for the learner, thus guiding him through the system towards goals he may take some small part in determining. Its second purpose is to make possible logical selection of appropriate materials that will offer the learner a variety of routes through the knowledge at a pace largely determined by himself. Just as the good teacher wishes to know his individual pupils well enough to take into account their peculiar needs, so in the computer-managed system a data bank of student characteristics, including details of recent performance on tests, is intended to form the basis for 'individually prescribed instruction' or for 'learning in accordance with needs'.

The first problems arise in the computer-based system when the data bank has to be filled. Consider the Open University, the simplest system of the three examples when its data bank of student characteristics is concerned. The bank contains, for each student, a minimum of biographical data collected on his application form. Few, if any, of the details are of value for those who construct the courses, and none at all are taken

into account to determine his course of study. Once he starts studying, his test results are also banked; some of this set of data does at least go to his tutor and counsellor, and may help those people to help him.

In PLAN and IPI, students provide large quantities of data through extensive testing carried out at various times but particularly at the end of the school year. For example, in PLAN, a student takes an annual test battery consisting of 1) an 18-scale Developed Abilities Performance Test, assessing such variables as vocabulary development, reading comprehension and abstract reasoning ability; 2) a 30-scale General Information Test to determine his pattern of functional interests; 3) a 12-scale PLAN Interest Inventory to ascertain areas of potential or expressed interest; 4) a Student Attitude Inventory; and 5) a battery of PLAN Achievement Tests. This battery is in addition to the tests he takes for each of the twenty or so modules he studies each year in at least four major school subjects. In PLAN, data are also collected from the teacher about the student's ability to work independently.

There are three kinds of problem that arise. They relate to the mode of data collection, the appropriateness of the tests, and the ways in which the data are used.

Firstly, the mode of data collection is open to criticism. Batteries of paper-and-pencil tests as used in PLAN seem the only practical way to collect data for the computer bank; lengthy structural interviews, observation of classroom behaviour and other psycho-sociological means of measuring are generally out of the question. Yet the batteries induce strong test-orientation in students (even adult Open University students) so that many regard the tests as the most important part of the programme, and only regard the learning materials as worthwhile if they relate closely to the tests. This is undesirable, since even large batteries supply only limited samplings of student attitudes and performance.

Secondly, the tests are often inappropriate, so that the data are of little use for diagnosis. Many of the test results leave the instructional designer saying 'so what?'. For example, what use can be made of an abstract reasoning score? If the score is low, can the student be asked to complete learning modules that will take the score into account, or produce an increase in it? The answer is no, because no modules have been designed with that particular construct in mind, nor do instructional designers have much idea of how to put together such modules. There are a few areas, such as mathematics, where the matching of test scores and learning materials is a little more straightforward, but even in those areas the decision-rules have not yet been devised for the learner to move from supposedly diagnostic scores to appropriate learning materials. Cooley and Glaser (1969) admit these difficulties in the IPI project.

Thirdly, the data easily get misused. Because the tests yield quantified data, these data quickly assume a validity and reliability far beyond the levels that would be quoted by the test designers. Since the data have to be used to guide difficult decisions about modules to be studied, the temptation to move into a mechanistic mode of thought is very strong. 'The data say you are a 2474 type, therefore you should study module 41-604-32.' This tendency is reinforced in PLAN, where the computer is supplied with a set of decision-rules relating chiefly to what the student already knows, plus his pace of study. The rules lead to the production by the computer of a programme of study for the student for the ensuing quarter or year. It is true that in PLAN (and at the Open University) the computer can be over-ruled by the teacher, but the computer appears to have far more diagnostic data at its fingertips. Its suggestions are usually followed, backed as they are by about 40 seconds of computer time, representing an extremely large number of decisions for each student. The rules themselves, however, are rather primitive: if a student has completed modules X and Y, he should now do Z, in so many days;

if his reading level is below a certain score, he should use the 'easy' unit (if one was prepared for that module); if he has a record of poor independent study habits, the teacher should be reminded of this for units requiring mostly independent study, and so on.

A more flexible but also more limited system of computer-based diagnostic assessment has been developed in a few small projects using on-line testing. For example, the University of California at Davis has given its students in animal and human physiology access to a computer terminal on which they can test their knowledge of course material and receive advice on remedial work (Walters, R.F. et al., 1972). This project is heavily dependent on tutor support on a ratio of 1:3 up to 1:12, therefore it cannot easily be compared with PLAN, IPI or the Open University.

Problems in the generation, analysis and banking of test-items

Stolurow (1965) was one of the first to use the computer to "invent" or generate test items. His interest lay in the direction of using the computer to help generate programmed learning sequences, particularly in mathematics. Several reports were published in the late 1960s by Stolurow and his colleagues at the University of Illinois, but these studies must be regarded as strictly experimental work.

Suppes, at Stanford University (see Suppes and Morningstar, 1970), has experimented with computer generation of test-items for mathematics and logic in his series of applied CAI studies. Searle, his chief educational assistant, indicated (in a personal interview in August 1973) that this had been one of the least successful aspects of their studies. She emphasised that to get the computer to the stage of generating test-items a great deal of preliminary pedagogical work was required. In the process of this work it was quite likely that the investigator would discover that the instructional materials were not rich enough to yield a large number of items. Computer generation of items depends on the existence of several

factors or variables that can be manipulated. For some subject-matters, Searle had found that only trivial variables were available for manipulation.

In Searle's opinion, it was very expensive to generate items by computer and seldom was the expense justified by the results. She thought there were relatively few situations in which such a large variety of items was required. Even the Open University's demand for a number (up to 15) of tests covering the same material, for use in different years, was probably too small to justify computer generation.

A further problem arose in connection with the computer-generated items in Suppes' project: in the schools, it proved difficult to persuade people who had not set the tests to accept the results for students they had taught. This was possibly a general problem of teacher attitudes towards an external agency, but the teachers felt that the tests were often inappropriate. The items may indeed have been a poor sampling of the curriculum taught by those teachers - even if the items were representative of some written curriculum. If the items were <u>not</u> representative, then the teachers would have a just complaint that the tests were apparently controlling the curriculum. As in many instructional situations, the relationships between curriculum, objectives and test items may not have been as good as they should have been.

Whether or not the computer is used to generate test items, it can certainly help to analyse them, both in terms of content and performance (Baker, 1971). Content analysis of items is carried out for indexing purposes, and the computer is used simply to provide a fast access system. Thus within an item bank there will be groupings of items by subject-matter, course level, question-type, and so on. A complete test can be put together according to any set of specifications using these groupings, once there are sufficient items in the bank. The problems of compilation originate in difficulties of content recognition (see Garbner, et al., 1969) and selecting suitable descriptors, difficulties not peculiar to computers (Committee on Information in the Behavioral Sciences, 1967).

The content specifications would be regarded as useless, however, without performance data. Many years of psychometric work have led to a set of standard requirements of data for items in a test-item bank. These requirements include difficulty and discriminate indices, reliability indices, and descriptions of the samples on which these indices are based. The discrimination index shows the extent to which students who do well on the test as a whole do well on individual items: i.e., does the item discriminate between good and poor performers? The reliability index reflects the consistency with which the item appears to test performances; for example, an ambiguously-worded item is likely to yield low reliability. The descriptions of samples used in deriving the indices are vital for the test compiler to be able to judge whether the items are appropriate for the group he intends to have tested.

There is no problem in providing these basic performance data through the computer, which is ideal for analysing the scores of students who have already taken the test items. This is standard practice in some North American universities, which make wide use of objective tests. In fact, the computer can quite easily provide a number of additional indices, useful to the test author in reviewing his items and in deciding which to eliminate or revise (Thorndike, 1971). A quite sophisticated system of this kind is being prepared for use in the Open University (for an outline of the proposal, see Lewis, 1972).

The essence of a computerised test-item bank is that the institution has not only a reserve store of items ready to be assembled into tests, at short notice if necessary, but also a store of items of known characteristics. Yet the following examples indicate that frequently this point is ignored.

At a junior college in Joliet, Illinois, a mini-bank of items was developed on a small computer to serve the needs of a biology for multiple versions of a test used in an audio-tutorial course. The report on this work (Wagner and Bleed, 1972) indicates that items were selected at random from the bank,

not on the basis of known characteristics. This is not really satisfactory. A similar bank at Washington State University lacked not only item characteristics but yielded astonishingly low reliability coefficients (Hammer and Henderson, 1972) because of poor psychometric design. The EXAMINER test-item bank used at the College of Business Courses at the University of South Florida is also a primitive tool, psychometrically speaking (Birkin, 1972). A bank developed at Iowa State University for a biology course did at least have difficulty estimates attached to each item, but these were not refined on the basis of student performance (Franke, et al., 1972).

The value of knowing the characteristics of the items, at least in terms of given test samples, is that tests for different purposes or parallel forms of a test for the same purpose, can be compiled. For example, many university tests compiled intuitively have a large proportion of medium-difficulty items, yet the aims of the examiners include isolating the excellent and the failing students. The difference in practice between the failing and passing students may lie in one vital item, failed by the failures and passed by the rest. Instead, it would be better to have items of that level of difficulty and fewer middle-difficulty items.

In summary, item banks have a disappointing record. The problems are not insuperable, but few institutions have tried to tackle them.

Problems in the sequencing of learning materials

Bruner (1960) suggested that schoolchildren of any age could be taught any concept, admittedly at different levels of understanding. Gagné (1965, 1968), on the other hand, argued for the hierarchical nature of learning tasks. In a recent seminar, Pask proposed that a nodal structure can be identified for any field of subject-matter, and that learners may enter that structure at many different points with fair chances of success in learning its contents.

Of these three authorities, Gagné is the one chiefly espoused by those working with computer-managed learning. This is not surprising, since he

is able to supply a more dogmatic set of decision-rules for sequencing learners' use of materials.

The basic problem in using the computer to determine the sequence each learner should follow is one of finding reasons for reducing the options open to the student at each decision point. For example, if a student in PLAN has just completed module 21-540-2, Multiplication and Division of Fractions, what is there to prevent him from going next to any module in PLAN? For a series of ten objectives, there are over 3 million possible sequences.

To begin with, there are practical restrictions. For example, the supply of instructional materials in any one classroom is limited by cost and space consideration. Third graders cannot study high school modules, and even wealthy American schools are forced to assume for economic reasons, that no more than a dozen students will arrive at the same module simultaneously: the cost of equipping classrooms would otherwise be very high indeed.

In addition, there are theoretical restrictions, chiefly in mathematics, when the sequence of learning tasks is thought to be critical. Even then, however, it is interesting to note that there was some disagreement among the mathematicians working on PLAN about the sequences to be followed. In the event, the chosen sequences are similar to those in the available text-books, because the student is less likely to become confused if he follows them!

For most of the modules, however, the sequencing rules are extremely arbitrary, and may actually detract from the individualisation of learning, since the student will have to ignore the computer's authoritarian instructions in his programme of study if he wants to change the sequence written into the computer for one that he likes better.

Dunn (1972) reported that in PLAN, a student's programme of studies would be determined by the computer taking into account his long range goals (vocational interests), state and local curriculum requirements, and his past achievements. His paper certainly does not explain clearly how this is done. He escapes being specific. From all that has been published on the use of long range goals in PLAN, there is no indication that these goals influence the sequence of studies except possibly at secondary shool level, where some subject areas may be deliberately excluded or included. State and local requirements have a similar effect: they may require that U.S. history shall be taught to all eighth-graders, for example. A student's past achievement scores are also unlikely to influence the sequence of studies, except by accelerating the student through subject-matter in which he has excelled and decelerating him in areas in which he has failed before.

In summary, the computer does little to resolve the perennial problems of sequencing faced by all curriculum reformers and designers. Decisions about whether a student learns set theory before he learns vulgar fractions remains in the hands of the textbook writers even in PLAN, since the modules are basically sets of instructions for using portions of existing texts. Certainly both PLAN and IPI are far from the stage of being able to use the computer to manage the sequence of learning for each student in a fully adaptive fashion.

Interface problems

People tend to blame machines. In computer-managed instruction, the computer comes in for more than its share of blame from administrators, teachers and students. Many of the problems can be traced to the interface, or the point at which the people actually have to use the machines by communicating with them. Interface problems arise mainly through incorrect signals being given to the computer.

These incorrect signals are of two kinds: incorrect programming of the computer and incorrect data. The programming of the computer is in the hands of the developer of the system, whether it is the Open University or American Institutes for Research/Westingnouse Learning Corporation (for PLAN). Errors in the development stage may cause problems when the system becomes operational, but these can be remedied in subsequent years.

The second type of error, incorrect data, is much harder to deal with, and deserves special attention here. The issue that is at stake is the validity of the data base used to reach decisions, immense as it is in each of the three systems, PLAN, IPI and the Open University. The problems to be solved are those of rendering the base more valid.

The first point at which learners meet the machine, as it were, is in entering the system. The applicant to the Open University receives a form designed to be data-processed. The limitations of the computer require him to code for the computer data about himself, about his occupation, for example. Most applicants are not accustomed to doing this. In PLAN, students taking the initial tests not only have to use an (ubiquitous) identification number; they also have to use IBM score-sheets that can be read by a mark-sensing terminal. These sheets not only require unaccustomed responses but also limit the types of questions that can be asked quite severely (27 five-choice items).

At the Open University, in PLAN and in IPI there are checking routines built into the programming to ensure that certain types of errors have not been made. Thus the document-readers will reject an incorrect identification number, will reject a score sheet on which a guessing student has marked more than a certain number of 'slots' for answers, and will reject a test for a module or course the student is not registered as studying.

None of the error-detection systems is entirely foolproof, but the estimates by the Open University computer administrators are that less than 5% of the documents enter the system with errors. This estimate may

be on the low side, however, since other documents may be filled in correctly from the computer's point of view but represent a range of errors on the part of the student in completing the document. For example, he may have skipped a line or two by mistake, thus throwing out the whole pattern of marks on the rest of the score-sheet. Computers can be programmed to search for this type of error, but at greater expense.

In PLAN, most schools have their own terminal, an IBM 2956 optical card reader plus an IBM 2740 control unit and teletype. Some of those involved in developing PLAN envisaged that students might have direct access to the terminal, but it soon became apparent that normally only teachers or clerks would be using it. During the night, the terminal is operated from a central computer (at one time, East and West Coast terminals were linked to an IBM 360/50 in Iowa), the cards are read, and the output for the teacher and his students is printed on the teletype, ready for the next day of school. The readers can only read 12 words a minute. These are too slow for a system that will have 60,000 students in 1973-74, but faster terminals are more expensive. Slow terminals mean not only high line costs but also delayed turnaround, which in turn may lead to reduced use of the computer facility. Indeed, Westinghouse is now selling PLAN as an individualised system with or without the computer-management. That decision was of course guided by commercial as well as educational motives.

In IPI, a PDP-15 computer is at the core of the system, being linked to IBM 1050 terminals and IBM 1232 scanners in the schools (Glaser, 1969).

Flanagan (1970) certainly pinned his hopes on PLAN students being able to use the computer to help them. He wrote 'It is possible that the main social impact of Project PLAN will be through its computerised guidance system by which it is hoped that the students will formulate reasonable goals, develop plans for achieving those goals, and take responsibility for carrying out those plans'. Yet, as we have seen, the articulation between the battery of tests used in PLAN and the instructional programme is weak.

The computer absorbs a mass of data, and by some devious magic, suggests a programme of studies. It is left to the teacher or counsellor to explain, if he can, the magic. This is one of the critical problems of interface. It can only be solved through establishing in the minds of teachers, counsellors and students the validity and reliability of the processes within the computer. There are those who will believe the computer printout simply because it is from the computer, but there are others who wish to have the processes explained. Or explanation and analysis, the processes reveal many defects. The fact that human planners and designers were responsible for these defects, and that human teachers and counsellors have many defects too, does not solve the interface problem.

In the Open University, similar problems of interface persist, chiefly in the computer-marked assignment system, but also in the survey research carried out by the University using computer-read questionnaires. Among students there is a basic distrust of the computer's ability to score the tests correctly; this is in spite of the fact that in every single case, students' appeals against faulty scoring by the computer have proved groundless. The students have not had great difficulty in learning how to use the forms, however, and accept the computer-marked assignments as part of the continuous assessment of their year's work. In the survey research, there have been the usual problems of students completing questionnaires for television programmes in weeks when there were none, and of misunderstanding directions for filling in the answer sheets for the document-reader. It is hard to assess whether these problems originate mainly in the format demanded by the reader (and hence the computer).

Conclusion

This paper has not listed every problem that has occurred or been thought of in attempts to implement computer-managed learning. Nor has it provided a taxonomy of such problems. Instead, it has discussed a selection of real problems that have to be recognised by those who wish to realise the computer's immense potential in helping to manage the learning process. The criticisms of the Open University, PLAN and IPI reflect upon areas where problems have not yet been solved, and ignore deliberately the great contributions that each of the three have made towards improving learning.

Finally, this paper is by no means a comprehensive review of the relevant literature. It is written from the viewpoint of an educational psychologist, not a computer specialist, and omits references to technical aspects of the computing or data processing equipment, for example. It also omits much relevant theoretical literature on curriculum design, testing and evaluation. Next time, perhaps a short book will be needed, rather than a short paper!

References

ANASTASIO, E. J. and MORGAN, J. S. (1972). Factors inhibiting the use of computers in instruction. Washington, D.C.: Interuniversity Communications Council.

BAKER, F. B. (1971). Automation of test scoring, reporting and analysis, in Thorndike, R.L. (ed.), Educational Measurement. Washington, D.C.: American Council on Education.

BIRKIN, S.J. (1972). An analysis of the use and effectiveness of EXAMINER - a computerised question bank and examination processing system - in College of Business courses at the University of South Florida, In Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

BRUNER, J. (1960). The process of education. Cambridge: Harvard University Press.

BUSHNELL, D.D. and ALLEN, D.W., eds. (1967). The computer in American education. New York: Wiley.

Committee on Information in the Behavioral Sciences (1967). Communications systems and resources in the behavioral sciences. Washington, D.C .: National Academy of Sciences.

COOLEY, W.W. and GLASER, R. (1968). An information and management system for individually prescribed instruction. Working Paper No.45. Pittsburgh: University of Pittsburgh Learning Research and Development Center.

COOLEY, W.W. and GLASER, R. (1969). The computer and individualised instruction. Science, 166, pp.574-82.

DUNN, J.A. (1972). The guidance program in the PLAN system of individualized education. Palo Alto: American Institutes for Research.

FLANAGAN, J.C. (1971). The PLAN educational system: a program for learning in accordance with needs. Palo Alto, Calif: American Institutes for Research.

FLANAGAN, J.C. (1970). The computer in Project PLAN. Palo Alto: American Institutes for Research.

FLANAGAN, J.C., DAILEY, J.T., SHAYCOFT, M.F., ORR, D.B. and GOLDBERG, I. (1962). The talents of American youth: Vol.I: Design for a study of American youth. Boston: Houghton Mifflin.

FRANKE, R.G., DOLPHIN, W.D., COVERT, G.F. and JORGENSEN, C.D. (1972). A computer-assisted method for teaching large enrolment lecture sections: the Biology Phase Achievement System, in Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

GAGNÉ, R.M. (1965). The conditions of learning. New York: Holt, Rinehart & Winston.

GAGNÉ, R.M. Learning hierarchies. Educ. Psychologist, Nov. 1968.

GEDDES, C.L. and KOOI, B.Y. An instructional management system for classroom teachers. Elementary School Journal, 69, 7, pp.337-45.

GERBNER, G., HOLSTI, O.R., KRIPPENDORFF, K., PAISLEY, W.J. and STONE, P.J. (eds.), (1969). The analysis of communication content: developments in scientific theories and computer techniques. New York, Wiley.

GLASER, R. (1969). Psychological questions in the development of computer-assisted instruction. Pittsburgh: University of Pittsburgh Learning Research and Development Center.

GLASER, R. and NITCO, A.J. (1970). Measurement in learning and instruction Pittsburgh: University of Pittsburgh Learning Research and Development Center.

HAMMER, M. and HENDERSON, C.O. (1972). Improving large enrollment undergraduate instruction with computer generated, repeatable tests, in Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

LATTA, R.F. and STRAUGHAN, J. (1972). Computer-managed instruction: a beginning and a reality at Western Washington State College, in Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

LEWIS, B.N. (1972). Course production at the Open University IV: The problem of assessment. British Journal of Educational Technology, 2, 3, pp.108-128.

PARKHURST, H. (1922). Education on the Dalton Plan. London: Bell.

STOLUROW, L.M. (1965). Computer-based instruction. Urbana: University of Illinois.

SUPPES, P and MORNINGSTAR, M. (1970). Four programs in computer-assisted instruction, in Holtzman, W.H. (ed.). Computer-assisted instruction testing and guidance. New York, Harper & Row.

THORNDIKE, R.L. (1971). Educational measurement for the seventies, in Thorndike, R.L. (ed.). Educational Measurement. Washington, D.C.: American Council on Education.

WAGNER, A and BLEED, R. (1972). Using a computer to support the testing program in audio-tutorial biology, in Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

WALTERS, R.F., NISHIMOTO, G.M., Horowitz, J.M. and BURGER, R.E. (1972). Initial development of individualised instruction with computer support, in Proceedings of the 1972 Conference on Computers in Undergraduate Curricula, June 12-14, 1972, Atlanta, Georgia. Atlanta: Southern Regional Education Board.

WASHBURNE, C., VOGEL, M. and GRAY, W.A. (1926). Results of practical experiments in fitting schools to individuals. Bloomington: Public School Publishing Co.

WEINGARTEN, F.W., NIELSEN, N.R., WHITELEY, J.R. and WEEG, G.P. (1973). A study of regional computer networks. Iowa City: University of Iowa.

Prepared by Professor David G. Hawkridge, Director of the Institute of Educational Technology at the Open University, Milton Keynes, Bucks., at the request of Mr. Richard Hooper, the Director of the National Development Programme in Computer-Assisted Learning.